arXiv:1404.4910v1 [cs.DC] 19 Apr 2014

Enumerating Maximal Bicliques from a Large Graph using MapReduce *

Arko Provo Mukherjee† Srikanta Tirthapura‡§

April 22, 2014

Abstract

We consider the enumeration of maximal bipartite cliques (bicliques) from a large graph, a task central to many practical data mining problems in social network analysis and bioinformatics. We present novel parallel algorithms for the MapReduce platform, and an experimental evaluation using Hadoop MapReduce.

Our algorithm is based on clustering the input graph into smaller sized subgraphs, followed by processing different subgraphs in parallel. Our algorithm uses two ideas that enable it to scale to large graphs: (1) the redundancy in work between different subgraph explorations is minimized through a careful pruning of the search space, and (2) the load on different reducers is balanced through the use of an appropriate total order among the vertices. Our evaluation shows that the algorithm scales to large graphs with millions of edges and tens of millions of maximal bicliques. To our knowledge, this is the first work on maximal biclique enumeration for graphs of this scale.

*A preliminary version of the paper "Enumerating Maximal Bicliques from a Large Graph using MapReduce" was accepted at the Proceedings of the 3rd IEEE International Congress on Big Data 2014.

†Department of Electrical and Computer Engineering, Iowa State University. Email: arko@iastate.edu

‡Department of Electrical and Computer Engineering, Iowa State University. Email: snt@iastate.edu

§The research reported in this papers is partially supported by the HPC equipment purchased through NSF MRI grant number CNS 1229081 and NSF CRI grant number 1205413.

1 Introduction

A graph is a natural abstraction to model rich relationships in data, and massive graphs are ubiquitous in applications such as online social networks [28, 33], information retrieval from the web [5], citation networks [4], and physical simulation and modeling [44], to name a few. Finding information from such data can often be reduced to a problem of mining features from massive graphs. We consider scalable methods for discovering densely connected subgraphs within a large graph. Mining dense substructures such as cliques, quasi-cliques, bicliques, quasi-bicliques etc. is an important area of study [1, 3, 17, 39].

A fundamental dense substructure of interest is a *biclique*. A biclique in a graph $G = (V, E)$ is a pair of subsets of vertices $L \subseteq V$ and $R \subseteq V$ such that (1) L and R are disjoint and (2) there is an edge $(u, v) \in E$ for every $u \in L$ and $v \in R$. For instance, consider the following graph relevant to an online social network, where there are two types of vertices, users and webpages. There is an edge between a user and every webpage that the user "likes" on the social network. A biclique in such a graph consists of a set of users U and a set of webpages W such that every user in U has liked every page in W. Such a biclique indicates a set of users who share a common interest, and is valuable for understanding the actions of users on this social network. Often, it is useful to identify only *maximal bicliques* in a graph, which are those bicliques that are not contained within any other larger bicliques. *We consider the problem of enumerating all maximal bicliques from a graph (henceforth referred to as MBE).*

Many graph mining tasks have relied on enumerating bicliques to identify significant substructures within the graph. For instance the analysis of web search queries [47] considered the "click-through" graph, where there are two types of vertices, web search queries and web pages. There is an edge from a search query to every page that a user has clicked in response to the search query. MBE was used in clustering queries using the click through graph. MBE has been used in social network analysis, in detection of communities in social networks [20], and in finding antagonistic communities in trust-distrust networks [25]. It has also been applied in detecting communities in the web graph [19, 35].

In bioinformatics, MBE has been used widely e.g. in construction of the phylogenetic tree of life [11, 31, 36, 46], structure discovery and analysis of protein-protein interaction networks [6, 37], analysis of gene-phenotype relationships [45], prediction of miRNA regulatory modules [48], modeling of hot spots at protein interfaces [22], and in analysis of relationships between

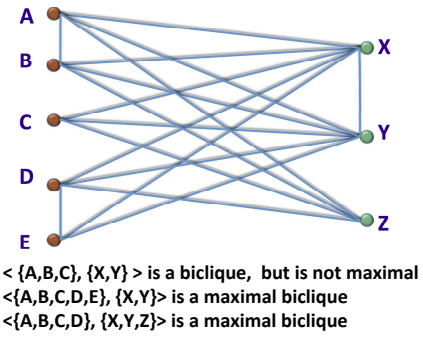

Figure 1: Maximal Bicliques

genotypes, lifestyles, and diseases [30]. Other applications include Learning Context Free Grammars [49], finding correlations in databases [18], for data compression [2], role mining in role based access control [7], and process operation scheduling [29].

While it is easy to find a single maximal biclique in a graph, enumerating all maximal bicliques is an NP-hard problem (Peters [34]). This does not however mean that typical cases are unsolvable. In fact, there are output-polynomial time algorithms whose theoretical runtime is bounded by a polynomial in the number of vertices in the graph, and the number of maximal bicliques that are output [3]. Thus it is reasonable to expect to be able to devise algorithms for MBE that work on large graphs, as long as the number of maximal bicliques output is not too high.

Current methods for enumerating bicliques have the following drawbacks. Most methods are sequential algorithms that are unable to make use of the power of multiple processors. For handling large graphs, it is imperative to have methods that can process a graph in parallel. Next, they have been evaluated only on small graphs of a few thousands of vertices and a few hundred thousand maximal bicliques, and have not been shown to scale

to large graphs. For instance, the popular "consensus" method for biclique enumeration [3] presents experimental data only on graphs of up to 2,000 vertices, and about 140,000 maximal bicliques, and other works [23, 24] are also similar. [1]. [2]

Our goal is to design a parallel method that can enumerate maximal bicliques in large graphs, with millions of edges and tens of millions of maximal bicliques, and which can scale with the number of processors.

1.1 Contributions

We present a parallel solution for MBE using the MapReduce framework [9]. At a high level, our approach clusters the input graph into overlapping subgraphs that can be processed independently in parallel, by different reducers. We implement the clustering approach using two different state-of-the-art sequential algorithms for MBE, one based on depth first search [24], and the other based on the consensus algorithm [3].

For this clustering approach to be effective on large graphs, we needed to augment it with two ideas that significantly improve the parallel performance. The first idea is concerned with reducing the overlap in the work done by different subtasks. It is usually not possible to assign disjoint subgraphs to different processors, and the subgraphs assigned to different tasks will overlap, sometimes significantly. Through a careful partitioning of the search space among the different tasks, we reduce redundant work among the tasks (this partitioning depends on details of the sequential algorithm used at each task).

The second idea is concerned with balancing the load between different tasks. With a graph analysis task such as biclique enumeration, the complexity of different subgraphs varies significantly, depending on the density of edges in the subgraph. Naively done, this can lead to a case where most reducers finish quickly, while only a few take a long time, leading to a poor parallel performance. We present a solution to keep the load more balanced, based on an ordering of vertices, which reduces enumeration load on subgraphs that are dense, and increases the load on subgraphs that are sparse, leading to a better parallel efficiency. We provide some basic statistical analysis of the Reducer runtimes with and without the load balancing to justify

[1] In our experiments, we show that the consensus and other sequential methods are unable to process our input graphs in a reasonable time

[2] It is not possible to quantify the complexity of a problem instance through the input size (number of vertices,and edges). However, the number of maximal bicliques, used in conjunction with the input size, is more indicative of the complexity.

our claim.

We give a detailed analysis of the communication costs of the clustering based MapReduce Algorithms described.

We present a direct parallelization of the consensus sequential algorithm [3], using an approach different from clustering. We found that while this approach may use a smaller memory per node that the clustering approach, it requires substantially greater runtime.

Finally, we present detailed experimental results on real-world and synthetic graphs. Overall, the clustering approach (using depth first search), when combined with load balancing and reduction of redundant work, performs the best on large graphs. Our algorithms can process graphs having millions of edges, and tens of millions of maximal bicliques, and can scale out with the cluster size. To our knowledge, these are the largest reported graph instances where bicliques have been successfully enumerated.

We also provide experimental evidence showing that our parallel Algorithm based on depth first search effectively generates only large maximal bicliques. For this we show that the runtime of our parallel algorithm decreases with increase in the size threshold of the generated maximal bicliques.

1.2 Prior and Related Work

Makino *et. al* [26] describes methods to enumerate all maximal bicliques in a *bipartite graph*, with the delay between outputting two bicliques bounded by a polynomial in the maximum degree of the graph. Zhang *et. al* [50] describe a branch-and-bound algorithm for the same problem. However, these approaches do not work for general graphs, as we consider here.

There is a variant of MBE where we only seek *induced* maximal bicliques in a graph. An induced maximal biclique is a maximal biclique which is also an induced subgraph; i.e. a maximal biclique $\langle L, R\rangle$ in graph G is an induced maximal biclique if L and R are themselves independent sets in G. We consider the *non-induced* version, where edges are allowed in the graph between two vertices that are both in L, or both in R (such edges are of course, not a part of the biclique). The set of maximal bicliques that we output will also contain the set of induced maximal bicliques, which can be obtained by post-processing the output of our algorithm. Note that for a bipartite graph, every maximal biclique is also an induced maximal biclique. Algorithms for Induced MBE include work by Eppstein [12], Dias *et. al* [10], and Gaspers *et. al* [14].

Alexe *et. al* [3] present an iterative algorithm for non-induced MBE using the "consensus" method, which we briefly review in Section 2.2. Another

technique for MBE is based on a recursive depth first search (DFS) [23, 24]. [23] presents an approach based on a connection with mining closed patterns in a transactional database, and apply the algorithm from [42], which is based on depth first search. [24] present a more direct algorithm for biclique enumeration based on depth first search, which we use in our work. This is described in more detail in Section 2.2.

Another approach to MBE is through a reduction to the problem of enumerating maximal cliques, as described by Gély *et. al* [15]. Given a graph G on which we need to enumerate maximal bicliques, a new graph G' is derived such that through enumerating maximal cliques in G' using an algorithm such as [40, 41], it is possible to derive the maximal bicliques in G. However, this approach is not practical for large graphs since in going from G to G', the number of edges in the graph increases significantly.

To our knowledge, the only prior work on parallel algorithms for MBE is by Nataraj and Selvan [32], who use the correspondence between maximal bicliques and closed patterns [23] to derive a parallel method for enumerating maximal bicliques. A significant issue is that [32] assumes that the input graph is presented as an adjacency matrix, which is then converted into a transactional database and distributed among the processors. In contrast, we do not assume an adjacency matrix, but assume that the graph is presented as a list of edges. Thus we are able to work on much larger graphs than [32]; the largest graph that they consider has 500 vertices and about 9000 edges.

MBE is related to, but different from the problem of finding the largest sized biclique within a graph (maximum biclique). There are a few variants of the maximum biclique problem, including maximum edge biclique, which seeks the biclique in the graph with the largest number of edges, and maximum vertex biclique, which seeks a biclique with the largest number of edges; for further details and variants, see Dawande *et al.* [8]. MBE is harder than finding a maximum biclique, since it enumerates all maximal bicliques, including all maximum bicliques.

2 Preliminaries

We present a formal problem definition, review prior sequential algorithms, and then briefly review the MapReduce parallel programming model that we use.

2.1 Problem Definition

We consider a simple undirected graph $G = (V, E)$ without self-loops or multiple edges, where V is the set of all vertices and E is the set of all edges of the graph. Let $n = |V|$ and $m = |E|$. Graph $H = (V_1, E_1)$ is said to be a sub-graph of graph $G = (V, E)$ if $V_1 \subset V$ and $E_1 \subset E$. H is known as an induced subgraph if E_1 consists of all edges of G that connect two vertices in V_1. For vertex $u \in V$, let $\eta(u)$ denote the vertices adjacent to u. For a set of vertices $U \subseteq V$, let $\eta(U) = \bigcup_{u \in U} \eta(u)$. For vertex $u \in V$ and $k > 0$, let $\eta^k(u)$ denote all vertices that can be reached from u in k hops. For $U \subseteq V$, let $\eta^k(U) = \bigcup_{u \in U} \eta^k(u)$. We call $\eta^k(U)$ as the k-neighborhood of U. For a set of vertices $U \subseteq V$, let $\Gamma(U) = \bigcap_{u \in U} \eta(u)$.

Definition 1. *A* biclique $B = \langle L, R \rangle$ *is a subgraph of* G *containing two non-empty and disjoint vertex sets,* L *and* R *such that for any two vertices* $u \in L$ *and* $v \in R$*, there is an edge* $(u, v) \in E$*.*

Note that the definition on $B = \langle L, R \rangle$ does not impose any restriction on the existence of edges among the vertices within R or within L, i.e., we consider *non-induced* bicliques.

Definition 2. *A biclique* $M = \langle L, R \rangle$ *in* G *is said to be a maximal biclique if there is no other biclique* $M' = \langle L', R' \rangle \neq \langle L, R \rangle$ *such that* $L \subset L'$ *and* $R \subset R'$*.*

The **Maximal Biclique Enumeration Problem (MBE)** is to enumerate the set of all maximal bicliques in graph $G = (V, E)$.

In our algorithms, we assume that vertex identifiers are unique and are chosen from a totally ordered set. This is usually not a limiting assumption, since vertex identifiers are usually strings which can be ordered using the lexicographic ordering.

2.2 Sequential Algorithms

We describe the two approaches to sequential algorithms for MBE that we consider, one based on a "consensus algorithm" [3], and the other based on depth first search [24].

Consensus Algorithm Alexe *et. al* [3] present an iterative approach to MBE. This type of algorithm starts off with a set of simple "seed" bicliques.

In each iteration, it performs a "consensus" operation, which involves performing a cross-product on the set of current candidates bicliques with the seed bicliques, to generate a new set of candidates, and the process continues until the set of candidates does not change anymore. Due to lack of space, we do not present the details here, and refer the reader to [3]. It is proved that these algorithms exactly enumerate the set of maximal bicliques in the input graph.

The consensus approach has a good theoretical guarantee, since its runtime depends on the number of maximal cliques that are output. In particular, the runtime of the MICA version of the algorithm is proved to be bounded by $O\left(n^3 \cdot N\right)$ where n is the number of vertices and N total number of maximal bicliques in G. The consensus algorithm has been found to be adequate for many applications and is quite popular.

We use the consensus algorithm in two ways. One as a candidate method for a sequential algorithm within each cluster. In another, we consider a direct parallelization of the consensus algorithm without using the clustering method.

Sequential DFS Algorithm The basic sequential depth first approach (DFS) that we use is described in Algorithm 1, based on [24]. It attempts to expand an existing maximal biclique into a larger one by including additional vertices that qualify, and declares a biclique as maximal if it cannot be expanded any further. The algorithm takes the following inputs: (1) the graph $G = (V, E)$, (2) the current vertex set being processed, X, (3) T, the tail vertices of X, i.e. all vertices that come after X in lexicographical ordering and (4) s, the minimum size threshold below which a maximal biclique is not enumerated. s can be set to 1 so as to enumerate all maximal bicliques in the input graph. However, we can set s to a larger value to enumerate only large maximal bicliques such that for $B =< L, R >$, we have $|L| \geq s$ and $|R| \geq s$. The size threshold s is provided as user input. The other inputs are initialized as follows: $X = \varnothing$, $T = V$.

The algorithm recursively searches for maximal bicliques. It increases the size of X by recursively adding vertices from the tail set T, and pruning away those vertices from T which along with X do not have any any common vertices in their neighborhood. From the expanded X, the algorithm outputs the maximal biclique $\langle \Gamma(\Gamma(X)), \Gamma(X) \rangle$.

Algorithm 1: Depth First Search: PA(G,X,T,s)

```
1  forall the vertex v ∈ T do
2      if |Γ(X ∪ {v})| < s then
3          T ← T \ {v}
4  if |X| + |T| < s then
5      return
6  Sort vertices in T as per ascending order of |η(X ∪ {v})|
7  forall the vertex v ∈ T do
8      T ← T \ {v}
9      if |X ∪ {v}| + |T| ≥ s then
10         N ← Γ(X ∪ {v})
11         Y ← Γ(N)
12         Biclique B ← ⟨Y, N⟩
13         if (Y \ (X ∪ {v})) ⊆ T then
14             if |Y| ≥ s then
15                 Emit B as a maximal biclique
16             PA(G, Y, T \ Y, s)
```

2.3 MapReduce

MapReduce [9] is a popular framework for processing large data sets on a cluster of commodity hardware. A MapReduce program is written through specifying map and reduce functions. The map function takes as input a key-value pair $\langle k, v\rangle$ and emits zero, one, or more new key-value pairs $\langle k', v'\rangle$. All tuples with the same value of the key are grouped together and passed to a reduce function, which processes a particular key k and all values that are associated with k, and outputs a final list of key-value pairs. The outputs of one MapReduce round can be the input to the next round. Communication happens only when the outputs from the map methods are retrieved by the different reduce methods based on the key, i.e. when data is grouped by keys. Further details are available in [9, 16]. We used Hadoop [38, 43], an open source implementation of MapReduce, on top of a distributed file system HDFS. While we consider the MapReduce framework for this work, our algorithms are generic and can be used with other distributed frameworks like Pregel [27].

3 MapReduce Algorithms for MBE

In this section, we describe algorithms for MBE using MapReduce. We first present the basic clustering approach, which can be used with any sequential algorithm for MBE, followed by enhancements to the basic clustering approach, and finally the parallel consensus approach.

3.1 Basic Clustering Approach

We first present the basic clustering framework for parallel MBE. The approach is to cluster the input graph into several overlapping sub-graphs (clusters) and then run the sequential DFS algorithm in parallel for each cluster.

For each $v \in V$, the cluster $C(v)$ consists of the induced subgraph on all vertices in $\eta^2(v)$ (i.e. the 2-neighborhood of v in G). The different clusters $C(v)$ are processed in parallel, and a sequential MBE algorithm is used to enumerate the maximal bicliques from each cluster. While all maximal bicliques in G are indeed output by this approach, the same biclique maybe enumerated multiple times. To suppress duplicates, the following strategy is used: a maximal biclique B arising from cluster $C(v)$ is emitted only if v is the smallest vertex in B according to the total order of the vertices. The basic clustering framework is generic and can be used with any sequential algorithm for MBE. We have considered the clustering algorithm using the DFS and the consensus algorithms for MBE.

Lemma 1. *The basic clustering approach enumerates all maximal bicliques in graph $G = (V, E)$.*

Proof. We show the following two properties. First, every maximal biclique in G must be output as a maximal biclique from cluster $C(v)$ for some $v \in V$. Second, every maximal biclique output from each cluster must be a maximal biclique in G. To prove the first direction, consider a maximal biclique $M = \langle L, R \rangle$ in G. Let v be the smallest vertex in M in lexicographic order, and without loss of generality suppose that $v \in L$. By the definition of a biclique, for each $u \in R$, u is a neighbor of v. Similarly, every vertex $w \in L$ is a neighbor of v, and is hence in $\eta^2(v)$. Hence M is completely contained in $C(v)$. Note that M is also a maximal biclique in $C(v)$. To see this, note that if M is not maximal biclique in $C(v)$, then M is not maximal in G either.

We prove by contradiction that every maximal biclique in each cluster $C(v)$ is also a maximal biclique in G. Consider a biclique M emitted as maximal from cluster $C(v)$ such that it is not maximal in G. Then, there exists

Table 1: Different versions of Parallel Algorithms based on Depth First Search (DFS)

Label	Algorithm
CDFS	Clustering based on Depth First Search (DFS)
CD0	CDFS + Reducing Redundant Work, without Load Balancing
CD1	CDFS + Reducing Redundant Work + Load Balancing using Degree
CD2	CDFS + Reducing Redundant Work + Load Balancing using Size of 2-neighborhood

a maximal biclique M' that can be generated by extending M. However, it is easy to see that every vertex in M' must also be contained in $\eta^2(v)$, and hence M' is also contained in $C(v)$, contradicting our assumption that M is a maximal biclique in $C(v)$. □

There are two problems with the basic clustering approach described above. First is *redundant work*. Although each maximal biclique in G is emitted only once, it may still be generated multiple times, in different clusters. This redundant work significantly adds to the runtime of the algorithm. Second is an *uneven distribution of load* among different subproblems. The load on subproblem $C(v)$ depends on two factors, the complexity of cluster $C(v)$ (i.e. the number and size of maximal bicliques within $C(v)$) and the position of v in the total order of the vertices. The earlier v appears in the total order, the greater is the likelihood that a maximum biclique in $C(v)$ has v has its smallest vertex, and hence the greater is the responsibility for emitting bicliques that are maximal within $C(v)$. Using a lexicographic ordering of the vertices may lead to a significantly increased workload for clusters of lower numbered vertices and a correspondingly low workload for clusters of higher numbered vertices.

3.2 Reducing Redundant Work

In order to reduce redundant work done at different clusters, we modify the sequential DFS algorithm for MBE that is executed at each reducer. We first observe that in cluster $C(v)$, the only maximal bicliques that matter are those with v as the smallest vertex; the remaining maximal bicliques in $C(v)$ will not be emitted by this reducer, and need not be searched for here. We use this to prune the search space of the sequential DFS algorithm used at the reducer.

All search paths in the algorithm which lead to a maximal biclique having a vertex less than v can be pruned away. Hence, before starting the DFS,

we prune away all vertices in the Tail set that are less than v, as described in Algorithm 6. Also, in DFS Algorithm 7, we prune the search path in Line 12 if the generated neighborhood contains a vertex less than v – maximal bicliques along this search path will not have v as the smallest vertex. Finally in Line 19 of Algorithm 7, we emit a maximal biclique only if the smallest vertex is the same as the key of the reducer in Algorithm 6.

The above algorithm, the "optimized DFS clustering algorithm", or "CD0" for short, is described in Algorithm 2. This takes two rounds of MapReduce. The first round, described in Algorithms 3 (map) and 4 (reduce), is responsible for generating the 1-neighborhood for each vertex. The second round, described in Algorithms 5 (map) and 6 (reduce) first constructs the clusters $C(v)$ and runs the optimized sequential pruning algorithm at the reducer. Note that Algorithm 6 passes the size threshold s while calling the optimized DFS Algorithm 7. The size threshold is an user input and can be passed on to Reducer (Algorithm 6) by using the *Configuration* parameters of Hadoop. Like the sequential algorithm, this parameter can be set to 1 to enumerate all maximal bicliques and to a larger value to enumerate only large maximal bicliques.

Algorithm 2: Algorithm CD0

Input: Edge List of $G = (V, E)$
1 Generate Adjacency List (Map) – Algorithm 3
2 Generate Adjacency List (Reduce) – Algorithm 4
3 Create Two Neighborhood (Map) – Algorithm 5
4 Create Two Neighborhood (Reduce) – DFS – Algorithm 6

Algorithm 3: Generate Adjacency List – Map

Input: Edge (x, y)
1 Emit ($key \leftarrow x$,$value \leftarrow y$)
2 Emit ($key \leftarrow y$,$value \leftarrow x$)

Lemma 2. *No maximal biclique in G is emitted by more than one reducer in Algorithm 7.*

Proof. Without the loss in generality, consider any maximal biclique $M = < L, R >$. Let $a \in \{L \cup R\}$ be the smallest vertex in $\{L \cup R\}$. Consider the reducer with $key = a$. In Line 20 of Algorithm 7, a maximal biclique is

Algorithm 4: Generate Adjacency List – Reduce

Input: $key \leftarrow v$,$value \leftarrow$\{Neighbors of v\}

1 $neighborhood \leftarrow \varnothing$
2 **forall the** $val \in value$ **do**
3 $\quad neighborhood \leftarrow neighborhood \cup val$
4 $N \leftarrow \langle v, neighborhood \rangle$
5 Emit ($key \leftarrow \varnothing$,$value \leftarrow N$)

Algorithm 5: Create Two Neighborhood – Map

Input: $N \leftarrow \langle v, neighborhood \rangle$

1 Emit ($key \leftarrow v$,$value \leftarrow N$)
2 **forall the** $y \in neighborhood$ **do**
3 $\quad$ Emit ($key \leftarrow y$,$value \leftarrow N$)

emitted only if the condition in line 18 is satisfied. This condition is satisfied by the reducer with $key = a$. However, this condition is not satisfied for any reducer such that $key \neq a$. Thus maximal biclique M is emitted only for the reducer with $key = a$. □

Lemma 3. *Algorithm 2 generates all maximal bicliques in a graph.*

Proof. The correctness of this Lemma can be proved from Lemmas 1 and 2. Algorithm 2 generates the 2-neighborhood induced sub–graph of each vertex in G. It then runs the Sequential DFS algorithm with the optimizations explained above.

Algorithm 6: Create Two Neighborhood (DFS) – Reduce

Input: $key \leftarrow v$, $value \leftarrow$ {2–neighborhood of v}

1 $G' = (V', E') \leftarrow$ Induced subgraph on $\eta^2(v)$
2 $X \leftarrow \varnothing$
3 $T \leftarrow V'$
4 **forall the** *vertex* $t \in T$ **do**
5 $\quad$ **if** $t < key$ **then**
6 $\quad\quad T \leftarrow T \setminus \{t\}$
7 CD0_Seq(G', X, T, key, s)

Algorithm 7: Optimized DFS – CD0_Seq

```
   Input: G′,X,T,key,s
 1 forall the vertex v ∈ T do
 2     if |Γ(X ∪ {v})| < s then
 3         T ← T \ {v}
 4 if |X| + |T| < s then
 5     return
 6 Sort vertices in T as per ascending order of |Γ(X ∪ {v})|
 7 forall the vertex v ∈ T do
 8     T ← T \ {v}
 9     if |X ∪ {v}| + |T| ≥ s then
10         N ← Γ(X ∪ {v})
11         Y ← Γ(N)
12         if Y contains vertices smaller than key then
13             continue
14         Biclique B ← ⟨Y, N⟩
15         if (Y \ (X ∪ {v})) ⊆ T then
16             if |Y| ≥ s then
17                 v_s ← Smallest vertex in B
18                 if v_s = key then
19                     // Maximal biclique found
20                     Emit (key ← ∅,value ← B)
21             CD0_Seq(G′, Y, T \ Y, key, s)
```

The correctness relies on the following two observations: Firstly, from Lemma 2, a maximal biclique is emitted from a reducer only if the smallest vertex in the biclique is same as the reducer key. Secondly, no vertex is ever removed from the set X. The set X thus always grows in size and never gets smaller in the course of the depth-first search. This is because the set Y is generated from set X in line 11 of Algorithm 7 and the set Y is passed as the new set X for the next level of recursion. The set Y is generated from the set X by taking the neighborhood of neighborhood of set X. $\eta(X)$ contains the set of all vertices connected to all vertices in X. Then $\Gamma(\Gamma(X))$ contains all vertices connected to all vertices in $\eta(X)$. This must include X. Hence $Y \supseteq X$.

From the above two observations we can prove the Lemma. Since we emit only those maximal bicliques for which the smallest vertices is the same as the reducer key k, we do not need to search the paths that produce maximal bicliques with smallest vertex less than k. Also, since no vertex is ever removed from the set X through the recursion path, we can be sure that at no point in the execution of the algorithm we will have $v \in X$ such that $v < k$. Now the set T can be considered as the candidate set as we always add elements to set X from set T. Thus in Algorithm 6 we remove all vertices from the set T that are less than k. Further in Algorithm 7, if we generate a maximal biclique in Line 12 with minimum vertex less than k then we prune the search tree through that path as all further maximal bicliques found in that search path will contain that vertex less than k. □

3.3 Load Balancing

Algorithm 8: Algorithms CD1 and CD2

Input: Edge List of $G = (V, E)$

1 Generate Adjacency List (Map) – Algorithm 3
2 Generate Adjacency List (Reduce) – Algorithm 4
3 Send vertex property (Map) – Algorithm 5
4 Send vertex property (Reduce) – Algorithm 9
5 Create Two Neighborhood (Map) – Algorithm 10
6 Create Two Neighborhood (Reduce) – Algorithm 11

Algorithm 9: Send vertex property – Reduce

Input: $key \leftarrow v$, $value \leftarrow$ {2–neighborhoods of v}

1 $S \leftarrow$ 2–neighbors of v
2 $neighborhood \leftarrow$ Compute neighborhood of v from S
3 $N \leftarrow \langle v, neighborhood \rangle$
4 `// Need to pass neighborhood for Round 3`
5 Emit ($key \leftarrow \varnothing$,$value \leftarrow N$)
6 `// Need to send vertex property to all 2--neighbors`
7 $p \leftarrow$ Value of vertex property of v from S
8 **forall the** *vertices* $s \in S$ **do**
9 ⌊ Emit($key \leftarrow \varnothing$,$value \leftarrow [s, v, p]$)

Algorithm 10: Create Two Neighborhood with vertex property – Map

Input: $N \leftarrow \langle v, neighborhood \rangle$ OR $[s, v, p]$
if $Input = N$ **then**
  Emit ($key \leftarrow v$,$value \leftarrow N$)
  forall the $y \in neighborhood$ **do**
    Emit ($key \leftarrow y$,$value \leftarrow N$)
else if $Value = [s, v, p]$ **then**
  Emit ($key \leftarrow s$,$value \leftarrow [v, p]$)

Algorithm 11: Create Two Neighborhood with vertex property (DFS) – Reduce

Input: $key \leftarrow v$, $value \leftarrow \{\eta^2(v)$ along with vertex properties$\}$
$G' = (V', E') \leftarrow$ Induced subgraph on $\eta^2(v)$
$Map \leftarrow$ HashMap of vertex and vertex property created from $value$ required to compute the new ordering
$X \leftarrow \varnothing$
$T \leftarrow V'$
forall the *vertex* $t \in T$ **do**
  if $t < key$ *in the new ordering* **then**
    $T \leftarrow T \setminus \{t\}$
CDL_Seq(G', X, T, key,Map, s)

In Algorithm 2, lexicographical ordering was used to order the vertices, which is agnostic of the properties of the cluster $C(v)$. The way the optimized DFS works (Algorithm 7), a reducer processing a vertex that is earlier in the total order is responsible for emitting more of the maximal bicliques within its cluster.

For improving load balance, we adjust the position of vertex v in the total order according to the properties of its cluster $C(v)$. Intuitively, the more complex cluster $C(v)$ is (i.e. more and larger the maximal bicliques), the higher should be position of v in the total order, so that the burden on the reducer handling $C(v)$ is reduced. While it is hard to compute (or estimate) the number of maximal bicliques in $C(v)$, we consider two properties of vertex v that are simpler to estimate, to determine the relative ordering of v in the total order: (1) Size of 1-neighborhood of v (Degree), and (2) Size of 2-neighborhood of v

Algorithm 12: Load Balanced DFS – CDL_Seq

```
   Input: G′,X,T,key,Map,s
 1 forall the vertex v ∈ T do
 2     if |Γ(X ∪ {v})| < s then
 3         T ← T \ {v}
 4 if |X| + |T| < s then
 5     return
 6 Sort vertices in T as per ascending order of |Γ(X ∪ {v})|
 7 forall the vertex v ∈ T do
 8     T ← T \ {v}
 9     if |X ∪ {v}| + |T| ≥ s then
10         N ← Γ(X ∪ {v})
11         Y ← Γ(N)
12         if Y contains vertices smaller than key in the new ordering
           then
13             continue
14         Biclique B ← ⟨Y, N⟩
15         if (Y \ (X ∪ {v})) ⊆ T then
16             if |Y| ≥ s then
17                 v_s ← Smallest vertex in B in the new ordering
18                 if v_s = key then
19                     // Maximal biclique found
20                     Emit (key ← ∅,value ← B)
21             CDL_Seq(G′, Y, T \ Y, key,Map, s)
```

In case of a tie, the vertex ID is used as a tiebreaker. These approaches were considered since vertices with higher degree are potentially part of a denser part of the graph and are contained within a greater number of maximal bicliques. The size of the 2-neighborhood gives the number of vertices in $C(v)$ and may provide a better estimate of the complexity of handling $C(v)$, but this is more expensive to compute than the size of the 1-neighborhood of the vertex.

The discussion below is generic and holds for both approaches to load balancing. To run the load balanced version of DFS, the reducer running the sequential algorithm must now have the following information for the vertex

(key of the reducer) : (1) 2-neighborhood induced subgraph, and (2) vertex property for every vertex in the 2-neighborhood induced subgraph, where "vertex property" is the property used to determine the total order, be it the degree of the vertex or the size of the 2-neighborhood. The second piece of information is required to compute the new vertex ordering. However, the reducer of the second round does not have this information for every vertex in $C(v)$, and a third round of MapReduce is needed to disseminate this information among all reducers. Further details are described in Algorithm 8. The DFS sequential algorithm for load balancing, described in Algorithm 11, is the same as the optimized DFS sequential algorithm 6, except that it orders using the vertex property (ties broken by IDs) rather than the simple lexicographic ordering.

3.4 Communication Complexity

We consider the communication complexity of Algorithms CD0, CD1 and CD2. For input graph $G = (V, E)$, we know that $n = |V|$ and $m = |E|$. Let us assume Δ to be the largest degree and $\bar{d}$ to be the average degree of vertices, where $\bar{d} = m/n$. Also, let β be the Output Size.

Definition 3. *Communication complexity of a MapReduce Algorithm $\mathcal{A}$ for Round r is denoted by $\mathcal{C}^r_{\mathcal{A}}$ and is defined as the sum of the total number of bytes emitted by all Mappers and the total number of bytes emitted by all the Reducers. We consider the output size for Reducers contributing to $\mathcal{C}^r_{\mathcal{A}}$ as each Reducer writes into the distributed file system incurring communication.*

Definition 4. *Let $\mathcal{C}_{\mathcal{A}}$ denote the total communication complexity for a MapReduce Algorithm $\mathcal{A}$ having R rounds. We define $\mathcal{C}_{\mathcal{A}} = \sum_{r=1}^{R} \mathcal{C}^r_{\mathcal{A}}$.*

Lemma 4. *Total communication complexity of Algorithm CD0 is $O(m \cdot \Delta + \beta)$.*

Proof. Algorithm CD0 has two rounds of MapReduce. For the first round, Algorithm 3, which is the Map method emits each edge twice, resulting in a communication complexity of $O(m)$. Similarly, Algorithm 4, which is the reducer emits each adjacency list once. This also results in a communication complexity of $O(m)$. Hence total communication complexity of the first round is $O(m)$.

Now let us consider the second round of MapReduce. The total communication between the Map and Reduce methods (Algorithms 5 and 6 respectively) can be computed by analyzing how much data is received by all

Reducers. Each reducer receives the adjacency list of all the neighbors of the key. Let d_i be the degree of vertex v_i, for $v_i \in V$, $i = 1, .., n$. Total communication is thus $\sum_{i=1}^{n} {d_i}^2 = {d_1}^2 + {d_2}^2 + \cdots {d_n}^2$. This is $O\left((d_1 + d_2 + \cdots + d_n) \cdot \Delta\right)$ $= O\left(n \cdot \bar{d} \cdot \Delta\right)$. Since $\bar{d} = m/n$, total communication becomes $O\left(m \cdot \Delta\right)$. The output from the final Reducer (Algorithm 6) is the collection of all maximal bicliques and hence the resulting communication cost is $O\left(\beta\right)$. Combining two rounds, total communication complexity becomes $O\left(m + m \cdot \Delta + \beta\right)$. $= O\left(m \cdot \Delta + \beta\right)$. □

Lemma 5. *Total communication complexity of Algorithm CD1 / CD2 is* $O\left(m \cdot \Delta + \beta\right)$.

Proof. First, note that both Algorithms CD1 and CD2 have the same communication complexity and observe that the first round uses the same Map and Reduce methods as CD0. Thus communication for Round 1 is $O\left(m\right)$. Again, note that Map method for Round 2 is same as CD0 and hence by Lemma 4, communication for Round 2 is $O\left(m \cdot \Delta\right)$.

The Reducer (Algorithm 9) of Round 2 sends the vertex property information to all its 2–neighbors. Thus every reducer receives information about all of its 2–neighbors. This makes the total output size of Reducer to be $O\left(m \cdot \Delta\right)$. The Map method of Round 3 (Algorithm 10) sends out the 2–neighborhood information as well as the vertex information to all vertices in 2–neighborhood. Thus communication cost becomes $O\left(m \cdot \Delta\right)$. The Reducer (Algorithm 11) emits all maximal bicliques and hence the resulting communication cost is $O\left(\beta\right)$. Thus total communication cost for Algorithms CD1 and CD2 is is $O\left(m \cdot \Delta + \beta\right)$. □

3.5 Parallel Consensus

We briefly describe another approach which directly parallelizes the consensus sequential algorithm of [3], in a manner different from the clustering approach. The motivation for this approach is as follows. The clustering approach has the following potential drawback, it requires each cluster $C(v)$ to have the entire 2-neighborhood of v. For dense graphs, the size of the 2-neighborhood of a vertex can be large, so that the complexity of each reduce task can be substantial. With the nature of the MapReduce model, dynamic load balancing among the reducers is not (easily) possible, so that load balancing will always be an issue for non-uniform, irregular computations.

Table 2: Various properties of the input graphs used, and runtime of different algorithms using 100 reducers. DNF means that the algorithm did not finish in 12 hours. The size threshold was set as 1 to enumerate all maximal bicliques. Runtime includes overhead of all MapReduce rounds including graph clustering, i.e. formation of 2–neighborhood.

Label	Input Graph	#vertices	#edges	#max–bicliques	Output Size	CDFS	CD0	CD1	CD2
1	p2p-Gnutella09	8114	26013	20332	407558	113	92	132	130
2	email-EuAll-0.6	125551	168087	292008	9161154	42023	4640	683	626
3	com-Amazon	334863	925872	706854	12739908	186	113	185	221
4	amazon0302	262111	1234877	886776	14553776	396	272	151	153
5	com-DBLP-0.6	251226	419573	1875185	82814962	1659	409	374	478
6	email-EuAll-0.4	175944	252075	2003426	111370926	DNF	DNF	6365	4154
7	ego-Facebook-0.6	3928	35397	6597716	315555360	8657	3858	1512	2943
8	loc-BrightKite-0.6	49142	171421	10075745	777419528	28585	11451	2506	2998
9	web-NotreDame-0.8	150615	300398	19941634	942300172	DNF	DNF	1688	2327
10	ca-GrQc-0.4	5021	17409	16133368	3101214314	37279	6895	5790	6374
11	ER-50K	50000	275659	51756	1116752	96	89	133	136
12	ER-60K	60000	330015	61821	1334716	98	89	135	135
13	ER-70K	70000	393410	71962	1589408	98	90	135	132
14	ER-80K	80000	448289	81983	1809070	102	90	136	134
15	ER-90K	90000	526943	92214	2125544	109	96	142	140
16	ER-100K	100000	600038	102663	2421528	114	97	144	143
17	ER-250K	250000	1562707	252996	6274864	167	114	165	162
18	ER-500K	500000	3751823	506319	15057870	374	167	251	252
19	Bipartite-50K-100K	150000	1999002	306874	9256056	873	183	227	253

Unlike the parallel DFS algorithm which works on subgraphs of G, the consensus algorithm is always directly dealing with bicliques within graph G. At a high level, it performs two operations repeatedly (1) a "consensus" operation, which creates new bicliques by considering the combination of existing bicliques, and (2) an "extension" operation, which extends existing bicliques to form new maximal bicliques. There is also a need for eliminating duplicates after each iteration, and also a step needed for detecting convergence, which happens when the set of maximal bicliques is stable and does not change further.

We developed a parallel version of each of these operations, by performing the consensus, extension, duplicate removal, and convergence test using MapReduce. We omit the details due to lack of space, but present experimental results from our implementations.

4 Experimental Results

We implemented the parallel algorithms on a Hadoop cluster, including the DFS-based clustering algorithms (CDFS, CD0, CD1, CD2), consensus-based clustering algorithm, and the parallel consensus algorithm. The Hadoop cluster has 28 nodes, each with a quad-core AMD Opteron processor with 8GB of RAM. All programs were written using Java and run with 2GB of

heap space. The Java version on all systems was 1.5.0, and the Hadoop version used was 0.20.203.

In addition, we implemented the sequential DFS algorithm [24] (without the optimizations that we introduced in our work), and the sequential consensus algorithm (MICA) [3]. The sequential algorithms were not implemented on top of Hadoop and hence had no associated Hadoop overhead in their runtime. But on the real-world graphs that we considered, these algorithms did not complete within 12 hours, except for the p2p-Gnutella09 graph.

We used both synthetic and real-world graphs for our experiments. A summary of all the graphs used is shown in Table 2. The real-world graphs were obtained from the SNAP collection of large networks [21] and were drawn from social networks, collaboration networks, communication networks, product co-purchasing networks, and internet peer-to-peer networks. Some of the real world networks were so large and dense that no algorithm was able to process them. For such graphs, we thinned them down by deleting edges with a certain probability. This makes the graphs less dense, yet preserves some of the structure of the real-world graph. We show the edge deletion probability in the name of the network. For example, graph "ca-GrQc-0.4" is obtained from "ca-GrQc" by deleting each edge with probability 0.4.

Synthetic graphs are either random graphs obtained by the Erdos-Renyi model [13], or random bipartite graphs obtained using a similar model. To generate a bipartite graph with n_1 and n_2 vertices in the two partitions, we randomly assign an edge between each vertex in the left partition to each vertex in the right partition. A random Erdos-Renyi graph on n vertices is named "ER-$< n >$", and a random bipartite graph with n_1 and n_2 vertices in the bipartitions is called "Bipartite-$< n_1 >$-$< n_2 >$".

We seek to answer the following questions from the experiments.

1. What is the relative performance of the different methods for MBE?
2. How do these methods scale with increasing number of reducers?
3. How does the runtime depend on the input size and the output size?

Figure 2 plots the runtime data for the algorithms mentioned in Table 2. All data used for these plots was generated with 100 reducers. *The runtime data given for the Parallel Algorithms include the time required to run all MapReduce rounds including time required to construct 2–neighborhood etc.*

Table 3: Mean and Standard Deviation computation of all 100 reducer runtimes for Algorithms CD0, CD1 and CD2. The analysis is done for the Reducer of the last MapReduce round as it performs the actual Depth First Search computation.

loc-BrightKite-0.6	CD0	CD1	CD2
Average	1005.94	631.92	625.14
Variance	3470135.82	256859.25	302764.97
Standard Deviation	1862.83	506.81	550.24
ego-Facebook-0.6	CD0	CD1	CD2
Average	479.37	380.32	422.7
Variance	473875.57.82	80146.79	273575.95
Standard Deviation	688.39	283.10	523.04

Impact of the Pruning Optimization. From Figure 2, we can see that the optimizations to basic DFS clustering through eliminating redundant work make a significant impact to the runtime for all input graphs. For instance, in Figure 2d, on input graph email–EuAll–0.6 CD0, which incorporates these optimizations, runs 9 times faster than CDFS, the basic clustering approach without reducing redundant work.

Impact of Load Balancing From Figure 2, we also observe that for graphs on which the algorithms do not finish very quickly (within 200 seconds), load balancing helps significantly. In Figure 2d, for graph email–EuAll–0.6, the Load Balancing approaches (CD1 and CD2) are 7 – 7.4 times faster than the Algorithm without Load Balancing, but optimized to reduce redundant work (CD0). Other examples include Figures 2e, where for input graph loc–BrightKite–0.6, CD1 was 4.5 times faster than CD0 and CD2 was about 3.8 times faster. For some graphs, such as email-EuAll-0.4 and web-NotreDame-0.8, CD0 failed to complete even after 12 hours, but CD1 and CD2 were successful in processing them within 2 hours. This shows that there must be significant imbalance in the load, with subproblems on some dense clusters being much more expensive to process than others. This shows that **for most input graphs, the versions optimized through both load balancing and reducing redundant work worked the best overall**.

However, for graphs that are quickly processed, the load balancing performs slightly slower than CD0 (see Figure 2a). This can be explained by

the additional overhead of load balancing (an extra round of MapReduce), which does not payoff unless the work done at the DFS step is significant.

We tried two Load Balancing approaches, one based on the vertex degree and the other on the size of the 2-neighborhood of the vertex. From Figure 2 we can observe no one approach was consistently better than the other, and the performance of the two were close to each other. For some input graphs, like Email-EuAll-0.4, the 2-neighborhood approach (CD2) fared better than the degree approach (CD1), whereas for some other input graphs like web-NotreDame-0.8, the degree approach fared better.

Finally, to further analyze the impact of Load Balancing, we also calculated the Mean and Standard Deviation of the run time of each of the 100 reducers of the Reducer Algorithm 6 for CD0 and Reducer Algorithm 11 for Algorithms CD1 and CD2 respectively. We present results of this analysis for input graphs loc-BrightKite-0.6 and ego-Facebook-0.6. Table 3 shows the mean and standard deviation analysis for the above-mentioned input graphs. **Observe that the Load Balanced Algorithms CD1 and CD2 have a much less standard deviation for reducer runtimes than CD0.**

Consensus versus Depth First Search Both clustering consensus as well as parallel consensus were much slower than the clustering DFS approaches. In all instances except for very small input graphs, clustering consensus was 6-11 times slower than CD1 and CD2 or worse (in many cases, clustering consensus did not finish within 12 hours while CD1 and CD2 finished within 1-2 hours). Further, direct parallel consensus, which uses a different parallelization strategy as explained in Section 3.5, was 13 to 100 times slower than clustering consensus. These numbers are not plotted in the figures. This shows that **Clustering based on DFS is the method of choice for parallel MBE.**

Scaling with Number of Reducers We measured how Algorithms CD1 and CD2 scaled with the number of reducers. In Figure 3 we plot the runtime of CD1 and CD2 with and increase in the number of reducers. In Figure 4, we also plot the speedup, defined as the ratio of the time taken with 1 reducer to the time taken with r reducers, as a function of the number of reducers r. We observe that the runtime decreases with increasing number of reducers, and further, the algorithms achieve near-linear speedup with the number of reducers. This data shows that the algorithms are scalable and may be used with larger clusters as well.

Relationship to Output Size: We also measured the runtime of the algorithms with respect to the output size. We define the output size of the problem as the sum of the number of edges for all enumerated maximal bicliques. Figure 5 shows the runtime of algorithms CD0, CD1, and CD2 as a function of the output size. This data is only constructed for random graphs, where the different graphs considered are generated using the same model, and hence have very similar structure. We observe that **the runtime increases almost linearly with the output size for all three algorithms CD0, CD1, and CD2**.

With real world graphs, this comparison does not seem as appropriate, since the different real worlds graphs have completely different structures; however, we observed that Algorithms CD1 and CD2 have a near-linear relationship with the output size, even on real world graphs.

Large Maximal Bicliques: Finally we measured the runtime of Algorithm CD1 with respect to the size threshold s. Figure 6 shows the runtime with respect to various size thresholds ranging from 1 to 5. **We can observe that the runtime decreases with increase in the size threshold s.** For instance, in Figure 6a, we can observe for the graph email-EuAll-0.4, the time taken to enumerate maximal bicliques larger than size 5, is about 5 times less and the time taken for size threshold 3 is about 2.1 times less than that to enumerate all maximal bicliques. Similarly, we can observe from Figure 6b, that enumeration times are 3.8 and 1.6 times less for size thresholds 5 and 3 respectively for graph web-NotreDame-0.8.

5 Conclusion

Mining maximal bicliques is a fundamental tool that is useful in uncovering dense relationships within data that can be represented as a graph. We presented scalable parallel methods for mining maximal bicliques from a large graph. We presented a basic clustering framework for parallel enumeration of bicliques. On top of this, we presented two optimizations, one for reducing redundant work, and the other for improving load balance, both of which significantly improved the observed performance. These algorithms scale well with increasing numbers of reducers. To our knowledge, our work is the first to successfully enumerate bicliques from graphs of this size; previous reported results were mostly sequential methods that worked on much smaller graphs.

The following directions are interesting for exploration (1) How does

this approach perform on even larger clusters, and consequently, larger input graphs? What are the bottlenecks here? and (2) Are there further improvements in search space pruning, and load balancing?

References

[1] James Abello, Mauricio G. C. Resende, and Sandra Sudarsky. Massive quasi-clique detection. In *LATIN 2002: Theoretical Informatics*, volume 2286 of *Lecture Notes in Computer Science*, pages 598–612. Springer Berlin Heidelberg, 2002. ISBN 978-3-540-43400-9.

[2] Pankaj K. Agarwal, Noga Alon, Boris Aronov, and Subhash Suri. Can visibility graphs be represented compactly? *Discrete & Computational Geometry*, 12(1):347–365, 1994.

[3] Gabriela Alexe, Sorin Alexe, Yves Crama, Stephan Foldes, Peter L. Hammer, and Bruno Simeone. Consensus algorithms for the generation of all maximal bicliques. *Discrete Applied Mathematics*, 145(1):11–21, 2004. URL `http://linkinghub.elsevier.com/retrieve/pii/S0166218X04000629`.

[4] Yuan An, Jeannette Janssen, and Evangelos E. Milios. Characterizing and mining the citation graph of the computer science literature. *Knowledge and Information Systems*, 6:664–678, 2004. ISSN 0219-1377. URL `http://dx.doi.org/10.1007/s10115-003-0128-3`. 10.1007/s10115-003-0128-3.

[5] Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. Graph structure in the web. *Computer Networks*, 33(1):309 – 320, 2000. ISSN 1389-1286.

[6] Dongbo Bu, Yi Zhao, Lun Cai, Hong Xue, Xiaopeng Zhu, Hongchao Lu, Jingfen Zhang, Shiwei Sun, Lunjiang Ling, Nan Zhang, Guojie Li, and Runsheng Chen. Topological structure analysis of the protein–protein interaction network in budding yeast. *Nucleic Acids Research*, 31(9): 2443–2450, 2003.

[7] Alessandro Colantonio, Roberto Di Pietro, Alberto Ocello, and Nino Vincenzo Verde. Taming role mining complexity in rbac.

Computers & Security, 29(5):548 – 564, 2010. ISSN 0167-4048. doi: http://dx.doi.org/10.1016/j.cose.2010.01.001. URL http://www.sciencedirect.com/science/article/pii/S0167404810000027.

[8] Milind Dawande, Pinar Keskinocak, Jayashankar M Swaminathan, and Sridhar Tayur. On bipartite and multipartite clique problems. *Journal of Algorithms*, 41(2):388 – 403, 2001. ISSN 0196-6774. doi: 10.1006/jagm.2001.1199. URL http://www.sciencedirect.com/science/article/pii/S019667740191199X.

[9] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51:107–113, January 2008. ISSN 0001-0782. doi: http://doi.acm.org/10.1145/1327452.1327492. URL http://doi.acm.org/10.1145/1327452.1327492.

[10] Vânia M. F. Dias, M.H. M. H. de Figueiredo Celina, and Jayme L. Szwarcfiter. Generating bicliques of a graph in lexicographic order. *Theoretical Computer Science*, 337:240–248, June 2005. ISSN 0304-3975.

[11] Amy C Driskell, Cécile Ané, J Gordon Burleigh, Michelle M McMahon, Brian C O'Meara, and Michael J Sanderson. Prospects for building the tree of life from large sequence databases. *Science*, 306(5699):1172–1174, 2004. doi: 10.1126/science.1102036. URL http://www.sciencemag.org/content/306/5699/1172.abstract.

[12] David Eppstein. Arboricity and bipartite subgraph listing algorithms. *Information Processing Letters*, 51:207–211, August 1994. ISSN 0020-0190.

[13] P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.

[14] Serge Gaspers, Dieter Kratsch, and Mathieu Liedloff. Graph-theoretic concepts in computer science. chapter On Independent Sets and Bicliques in Graphs, pages 171–182. Springer, 2008. ISBN 978-3-540-92247-6.

[15] Alain Gély, Lhouari Nourine, and Bachir Sadi. Enumeration aspects of maximal cliques and bicliques. *Discrete Applied Mathematics*, 157(7):1447 – 1459, 2009.

ISSN 0166-218X. doi: 10.1016/j.dam.2008.10.010. URL http://www.sciencedirect.com/science/article/pii/S0166218X08004563.

[16] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The Google File System. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, SOSP '03, pages 29–43. ACM, 2003. ISBN 1-58113-757-5. doi: http://doi.acm.org/10.1145/945445.945450. URL http://doi.acm.org/10.1145/945445.945450.

[17] David Gibson, Ravi Kumar, and Andrew Tomkins. Discovering large dense subgraphs in massive graphs. PVLDB, pages 721–732. VLDB Endowment, 2005. ISBN 1-59593-154-6.

[18] Christopher Jermaine. Finding the most interesting correlations in a database: how hard can it be? *Information Systems*, 30(1):21 – 46, 2005. ISSN 0306-4379. doi: http://dx.doi.org/10.1016/j.is.2003.08.004. URL http://www.sciencedirect.com/science/article/pii/S0306437903000802.

[19] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. Trawling the web for emerging cyber-communities. *Computer networks*, 31(11):1481–1493, 1999.

[20] Sune Lehmann, Martin Schwartz, and Lars Kai Hansen. Biclique communities. *Physical Review E*, 78:016108, Jul 2008. doi: 10.1103/PhysRevE.78.016108. URL http://link.aps.org/doi/10.1103/PhysRevE.78.016108.

[21] J. Leskovec. Stanford large network dataset collection. http://snap.stanford.edu/data/index.html.

[22] Jinyan Li and Qian Liu. 'double water exclusion' : a hypothesis refining the o–ring theory for the hot spots at protein interfaces. *Bioinformatics*, 25(6):743–750, 2009.

[23] Jinyan Li, Guimei Liu, Haiquan Li, and Limsoon Wong. Maximal biclique subgraphs and closed pattern pairs of the adjacency matrix: A one-to-one correspondence and mining algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 19(12):1625–1637, 2007.

[24] Guimei Liu, Kelvin Sim, and Jinyan Li. Efficient mining of large maximal bicliques. In *Data Warehousing and Knowledge Discovery*, volume 4081 of *Lecture Notes in Computer Science*, pages 437–448. Springer, 2006.

[25] David Lo, Didi Surian, Kuan Zhang, and Ee-Peng Lim. Mining direct antagonistic communities in explicit trust networks. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management*, CIKM '11, pages 1013–1018. ACM, 2011. ISBN 978-1-4503-0717-8. doi: http://doi.acm.org/10.1145/2063576.2063722. URL `http://doi.acm.org/10.1145/2063576.2063722`.

[26] Kazuhisa Makino and Takeaki Uno. New algorithms for enumerating all maximal cliques. In *Algorithm Theory-SWAT 2004*, pages 260–272. Springer, 2004.

[27] Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, SIGMOD '10, pages 135–146. ACM, 2010. ISBN 978-1-4503-0032-2. doi: 10.1145/1807167.1807184. URL `http://doi.acm.org/10.1145/1807167.1807184`.

[28] Alan Mislove, Massimiliano Marcon, Krishna P. Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, IMC '07, pages 29–42. ACM, 2007. ISBN 978-1-59593-908-1. doi: 10.1145/1298306.1298311. URL `http://doi.acm.org/10.1145/1298306.1298311`.

[29] Sylvain Mouret, Ignacio E. Grossmann, and Pierre Pestiaux. Time representations and mathematical models for process scheduling problems. *Computers & Chemical Engineering*, 35(6):1038 – 1063, 2011. ISSN 0098-1354. doi: http://dx.doi.org/10.1016/j.compchemeng.2010.07.007.

[30] R. A. Mushlin, A. Kershenbaum, S. T. Gallagher, and T. R. Rebbeck. A graph-theoretical approach for pattern discovery in epidemiological research. *IBM Systems Journal*, 46(1):135–149, 2007. ISSN 0018-8670. doi: 10.1147/sj.461.0135.

[31] N. Nagarajan and C. Kingsford. Uncovering genomic reassortments among influenza strains by enumerating maximal bicliques. In *IEEE International Conference on Bioinformatics and Biomedicine, 2008*, pages 223–230, 2008.

[32] R. V. Nataraj and S. Selvan. Parallel mining of large maximal bicliques using order preserving generators. *International Journal of Computing*, 8(3):105–113, 2009. ISSN 1727-6209.

[33] M E J Newman, D J Watts, and S H Strogatz. Random graph models of social networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl 1):2566–2572, 2002.

[34] René Peeters. The maximum edge biclique problem is np-complete. *Discrete Applied Mathematics*, 131:651–654, September 2003. ISSN 0166-218X.

[35] Jayson E. Rome and Robert M. Haralick. Towards a formal concept analysis approach to exploring communities on the world wide web. In *Formal Concept Analysis*, volume 3403 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2005. ISBN 978-3-540-24525-4.

[36] Michael J Sanderson, Amy C Driskell, Richard H Ree, Oliver Eulenstein, and Sasha Langley. Obtaining maximal concatenated phylogenetic data sets from large sequence databases. *Molecular Biology and Evolution*, 20(7):1036–1042, 2003. URL `http://www.ncbi.nlm.nih.gov/pubmed/12777519`.

[37] Regev Schweiger, Michal Linial, and Nathan Linial. Generative probabilistic models for protein–protein interaction networks–the biclique perspective. *Bioinformatics*, 27(13):i142–i148, 2011.

[38] K. Shvachko, H. Kuang, S. Radia, and R. Chansler. The hadoop distributed file system. In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1 –10, may 2010. doi: 10.1109/MSST.2010.5496972.

[39] Kelvin Sim, Jinyan Li, Vivekanand Gopalkrishnan, and Guimei Liu. Mining maximal quasi-bicliques to co-cluster stocks and financial ratios for value investment. In *Proceedings of the Sixth International Conference on Data Mining*, ICDM '06, pages 1059–1063. IEEE, 2006. ISBN 0-7695-2701-9. doi: 10.1109/ICDM.2006.111. URL `http://dx.doi.org/10.1109/ICDM.2006.111`.

[40] Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science*, 363:28–42, Oc-

tober 2006. ISSN 0304-3975. doi: 10.1016/j.tcs.2006.06.015. URL http://dl.acm.org/citation.cfm?id=1217595.1217600.

[41] Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. A new algorithm for generating all the maximal independent sets. *SIAM Journal on Computing*, 6(3):505–517, 1977.

[42] Takeaki Uno, Masashi Kiyomi, and Hiroki Arimura. Lcm ver.2: Efficient mining algorithms for frequent/closed/maximal itemsets. In *IEEE International Conference Data Mining Workshop Frequent Itenset Miing Implementations (FIMI)*. IEEE, 2004.

[43] Tom White. *Hadoop: The Definitive Guide*. O'Reilly Media, Inc., 1st edition, 2009. ISBN 0596521979, 9780596521974.

[44] Olga Wodo, Srikanta Tirthapura, Sumit Chaudhary, and Baskar Ganapathysubramanian. A graph-based formulation for computational characterization of bulk heterojunction morphology. *Organic Electronics*, 13(6):1105 – 1113, 2012. ISSN 1566-1199. doi: 10.1016/j.orgel.2012.03.007. URL http://www.sciencedirect.com/science/article/pii/S1566119912001036.

[45] Yang Xiang, Philip R. O. Payne, and Kun Huang. Transactional database transformation and its application in prioritizing human disease genes. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 9(1):294–304, 2012. ISSN 1545-5963. doi: 10.1109/TCBB.2011.58. URL http://dx.doi.org/10.1109/TCBB.2011.58.

[46] Changhui Yan, J Gordon Burleigh, and Oliver Eulenstein. Identifying optimal incomplete phylogenetic data sets from sequence databases. *Molecular Phylogenetics and Evolution*, 35(3):528–535, 2005. URL http://www.ncbi.nlm.nih.gov/pubmed/15878123.

[47] Jeonghee Yi and Farzin Maghoul. Query clustering using click-through graph. In *Proceedings of the 18th international conference on World wide web*, WWW '09, pages 1055–1056, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-487-4. doi: 10.1145/1526709.1526853. URL http://doi.acm.org/10.1145/1526709.1526853.

[48] Sungroh Yoon and Giovanni D. Micheli. Prediction of regulatory modules comprising micrornas and target genes. *Bioinformatics*, 21(2):ii93–ii100, 2005.

[49] Ryo Yoshinaka. Towards dual approaches for learning context-free grammars based on syntactic concept lattices. In *Developments in Language Theory*, volume 6795 of *Lecture Notes in Computer Science*, pages 429–440. Springer Berlin Heidelberg, 2011. ISBN 978-3-642-22320-4.

[50] Yun Zhang, Elissa J. Chesler, and Michael A. Langston. On finding bicliques in bipartite graphs: a novel algorithm with application to the integration of diverse biological data types. *Hawaii International Conference on System Sciences*, page 473, 2008. ISSN 1530-1605.

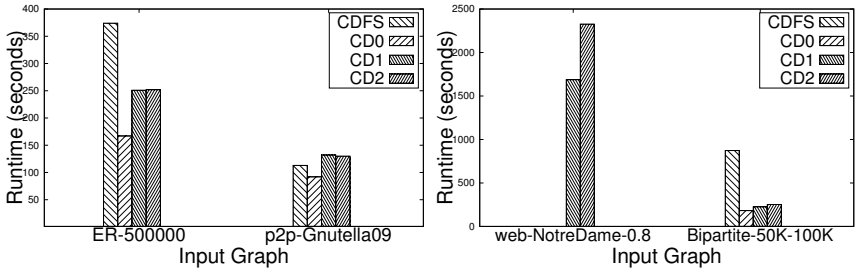

(a) ER-500K and p2p-Gnutella09 (b) web-NotreDame-0.8 and Bipartite-50K-100K

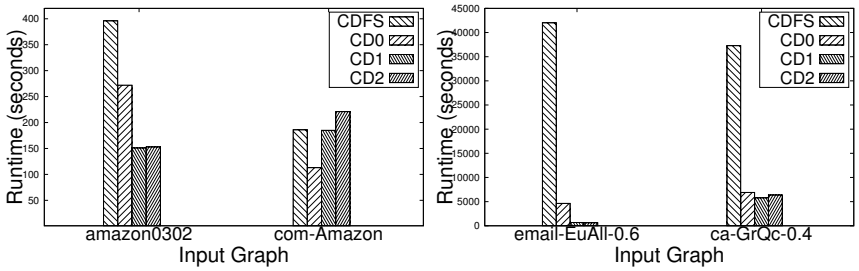

(c) amazon0302 and com-Amazon (d) email-EuAll-0.6 and ca-GrQc-0.4

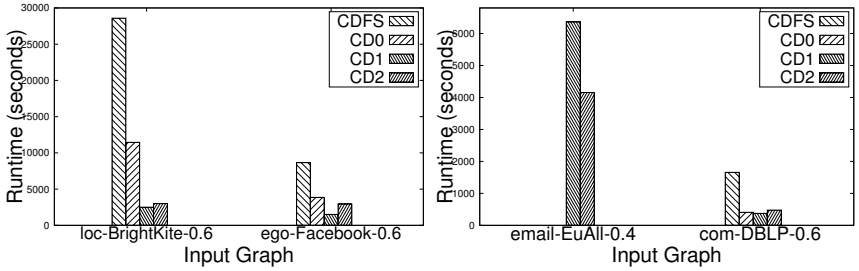

(e) loc-BrightKite-0.6 and ego-Facebook-0.6 (f) email-EuAll-0.4 and com-DBLP-0.6

Figure 2: Runtime of parallel algorithms on real and random graphs. If an algorithm failed to complete in 12 hours the result is not shown. All algorithms were run using 100 reducers. Runtime includes overhead of all MapReduce rounds including graph clustering, i.e. formation of 2–neighborhood.

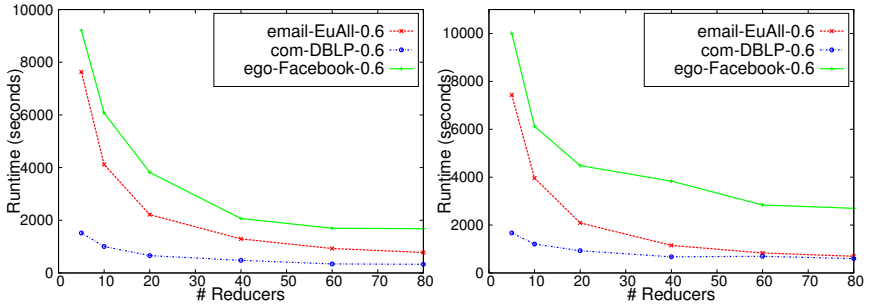

(a) Algorithm CD1 (b) Algorithm CD2

Figure 3: Runtime versus Number of Reducers.

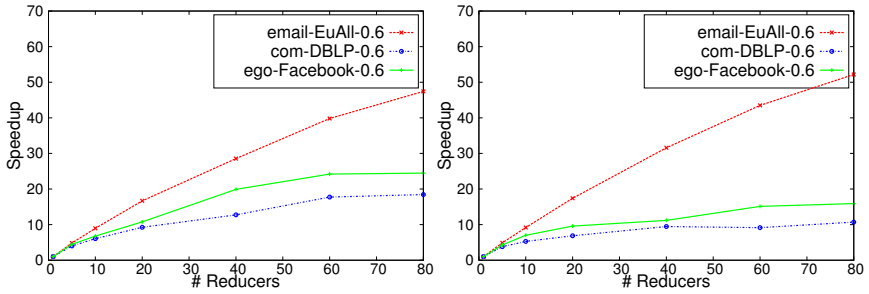

(a) Algorithm CD1 (b) Algorithm CD2

Figure 4: Speedup versus Number of Reducers.

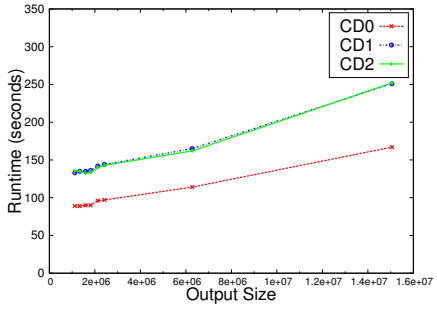

Figure 5: Runtime versus Output Size for random graphs. All Erdos-Renyi random graphs were used. Output size is defined as the number of edges summed over all maximal bicliques enumerated.

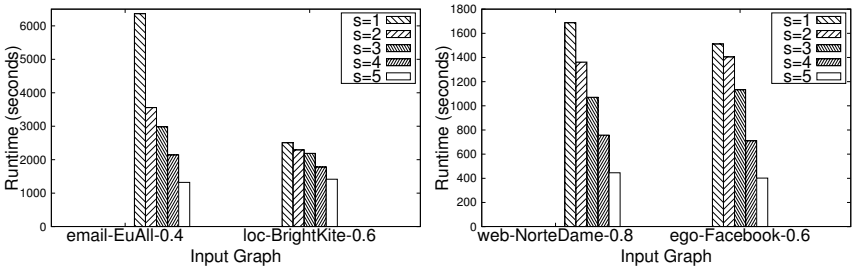

(a) email-EuAll-0.4 and loc-BrightKite-0.6

(b) web-NotreDame-0.8 and ego-Facebook-0.6

Figure 6: Runtime vs the size threshold for the emitted maximal bicliques. All experiments were performed using Algorithm CD1 and with 100 reducers.